White Paper Report

Report ID: 99026

Application Number: HD5105210

Project Director: David Koller (koller@virginia.edu)

Institution: University of Virginia

Reporting Period: 4/1/2010-3/31/2011

Report Due: 6/30/2011

Date Submitted: 11/1/2011

White Paper

NEH Digital Start-up Grant HD-51052-10

Project Title: Supercomputing for Digitized 3D Models of Cultural Heritage

Project Director: David Koller, University of Virginia

June 30, 2011

Supercomputing for Digitized 3D Models of Cultural Heritage

Background: High Performance Computing in the Humanities

The value of using high performance computing (HPC) for scientific applications is well-known. Several national supercomputing centers support scientific and engineering research, and most scientists at universities have access to local supercomputing resources. However, scholars working in the domain of "digital humanities" research have only more recently begun to explore the application of high performance computing to their fields of study. Examples include:

- The Institute for Computing in the Humanities, Arts, and Social Science (I-CHASS) at the University of Illinois at Urbana-Champaign supports a number of HPC humanities projects. For example, their recent "Culturomics 2.0" project used the large shared memory supercomputer Nautilus to analyze 100 million historical newspaper articles and perform data mining and analytics experiments.[1]
- The Perseus Digital Library at Tufts University is using advanced computational linguistic technologies to experiment with the analysis of ancient texts for the study of classics.
- In 2006, scholars at University College London studied the feasibility of using HPC technologies for analyzing historical census datasets.[2]
- Lev Manovitch and his cultural analytics team from the Software Studies Initiative at the University of California at San Diego are using Department of Energy supercomputers to analyze and visualize patterns in cultural images such as magazine covers and comic books.[3]
- A number of researchers are using supercomputers for computer graphic simulations and rendering. The University of Southampton archaeological visualization team uses the fastest academic supercomputer in the UK to interpret archaeological investigations at Catalhoyuk,[4] and the Scientific Computing and Imaging Institute used a 264-core supercomputer to render gigapixel multiview images of Michelangelo's David statue from the 3D scanned data of the Stanford Digital Michelangelo Project.[5]

---

[1] Nicole Hemsoth, "Nautilus Harnessed for Humanities Research, Future Prediction", *HPCwire*, Sept. 9, 2011**.**

[2] Melissa Terras, "The Potential and Problems in using High Performance Computing in the Arts and Humanities: the Researching e-Science Analysis of Census Holdings (ReACH) Project", *Digital Humanities Quarterly* (Fall 2009: v3 n4).

[3] Linda Vu, "A Computational Science Approach for Analyzing Culture", Berkeley Lab Computing Sciences, Feb.18, 2010.

[4] Graeme Earl et al., "High Performance Computing and the Construction of Archaeological Computer Graphic Simulations", XXXXI *Conference on Applications and Quantitative Methods in Archeology* (CAA 2011), April, 2011.

[5] "One billion polygons to billions of pixels", http://www.sci.utah.edu/news/60/431-visus.html.

Background: High Performance Computing in 3D Cultural Heritage

Of particular interest to us in this research project is the use of high performance computing for processing and analysis of digitized 3D cultural heritage objects. Three-dimensional (3D) digital data capture technologies, such as laser scanning, hold great promise for preserving and studying cultural heritage objects. Scanners allow the 3D surface shape of artifacts to be digitally measured with great precision, and humanities scholars have begun to utilize these technologies for building 3D digital models of archaeological artifacts and historical sites.

Cultural heritage scanning applications, however, can be particularly demanding relative to typical industrial data capture applications. Because of the emphasis on documentation and preservation for 3D scanning of cultural heritage artifacts, objects are usually scanned at uniform high sampling resolutions, which generates datasets of immense size. For example, the statue of Michelangelo's David digitized as part of the Stanford Digital Michelangelo Project was scanned at 0.25 millimeter resolution, in order to fully capture the chisel marks and marble surface wear, resulting in a total of over one billion point measurements across the full statue's surface. The modern time-of-flight scanners used to digitize large objects such as buildings are capable of gathering over 1,000,000 points per second of scanning, so digitizing an archaeological site or monument from several angles at full resolution can quickly yield huge datasets with billions of individual 3D measurements.

Unfortunately, traditional single-processor desktop computing environments have proved insufficient for processing and analyzing the huge cultural heritage datasets that are being created with 3D scanning technologies, as a single desktop computer has not possessed enough memory or CPU capability to efficiently handle many such raw data collections. In an effort to overcome these limitations, we are investigating the utilization of high performance parallel computing resources for processing and studying digitized 3D models of cultural heritage.

Project History

In 2007, we participated in the NEH Digital Humanities Initiative Workshop on Supercomputing & the Humanities, which discussed possible programs to foster collaboration between computer scientists and humanities scholars. In 2008, the NEH and Department of Energy (DOE) created the Humanities High Performance Computing Program (HHPC), an initiative to give humanities researchers access to powerful supercomputers at the DOE's National Energy Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory. Our project was selected to participate in this program for 2009, and was awarded up to one million compute hours on the NERSC computers.

This Start-Up grant was sought in order to supplement and support our efforts that have been initiated in part with this earlier HHPC grant. While the HHPC program provides access to considerable high performance computing resources, there was no direct support for personnel activities. We found that substantial time and effort was required to adapt software to run in parallel supercomputing environments, and to prepare input data and process the results. Thus the project funds primarily covered the effort of the Project Director to develop software, process data, and conduct the necessary research into applying supercomputing to the cultural heritage problems described below.

Project Activities

1. High performance computing for processing raw 3D scan data into completed 3D models. The raw 3D point cloud measurements output by 3D scanners require significant further processing steps to produce final, complete, watertight 3D surface models. These steps include aligning multiple scans from different positions into a single unified coordinate system ("registration"), merging the multiple aligned scans into a single consensus surface ("surface reconstruction"), and then filling any remaining small holes in the surface ("hole filling"). The geometric algorithms that perform these processing steps are very computationally demanding, and can require high performance computer resources to run efficiently on large datasets.

In the course of our project, we have adapted two particular existing algorithm implementations to run on massively parallel supercomputers. For registration, we adapted the non-rigid alignment method of Brown & Rusinkiewicz (2007).[6] For the surface reconstruction step, we extended the volumetric method of Curless & Levoy (1996).[7] In each case, there existed open-source software implementations of these algorithms, although further software development work was necessary to adapt them to run in parallel on our particular clusters and specialized supercomputing environments. Adapting the reconstruction software to run in parallel was relatively easy, as no shared-memory computation was necessary. No message passing API was required, as the parallel computation tasks simply shared their sub-results via the file system, and these sub-results were then combined to yield the final 3D model. The parallel algorithms were developed, tested, and run on clusters at the University of Virginia, Stanford University, and on massively parallel supercomputers at NERSC. The NERSC machines used included 'Jacquard', an Opteron cluster with 712 CPUs running a Linux operating system, and an IBM iDataPlex system with 3,200 cores. Although our codes can scale to any degree of parallelism, we rarely use in excess of 128 individual nodes.

We have exercised our high performance computing scan data processing pipeline with a variety of large cultural heritage 3D raw scan datasets. One example of notable statuary that we have processed is the Laocoön statue group from the Vatican Museums (Figs. 1 & 2), originally digitized in 2008. The final statue model that we produced from the 500 individual scan range images was at 0.5 millimeter resolution. The processing time required for the non-rigid alignment step was 2.1 hours utilizing 128 computation nodes, and the volumetric merging step to produce the surface model required 30 minutes while utilizing 64 nodes. Thus, an entire statue of this large size and fine resolution can be processed from raw data into a final surface model in well under 3 hours. Another example of a model that we both scanned and processed during our investigations for this project is a lead bust of Sir Walter Raleigh, from atop the Raleigh Tavern at Colonial Williamsburg (Fig. 3, left). The bust was scanned at 0.2 millimeter resolution with

---

[6] B. Brown and S. Rusinkiewicz, "Global Non-Rigid Alignment of 3-D Scans," *ACM Transactions on Graphics*, August 2007.

[7] B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images," *Proc. SIGGRAPH 1996*.

an articulated arm scanner, but processing at full resolution yielded a poor result due to miscalibration in one or more of the scanner joints. This problem could be rectified with an auto-calibration process, algorithms for which are well-suited to parallelization and HPC, but we have not yet attempted this. In addition to statuary, we have also have experimented with the processing of large architectural models, such as the 2.0 millimeter time-of-flight laser scans of the Charlton's Coffeehouse stone foundations at Colonial Williamsburg (Fig. 3, right) that we have previously captured. This dataset contains over 2 billion point measurements.

Figure 1: 3D scanning the Laocoön statue group

Figure 2: Raw scan data processed into 3D models of the Laocoön

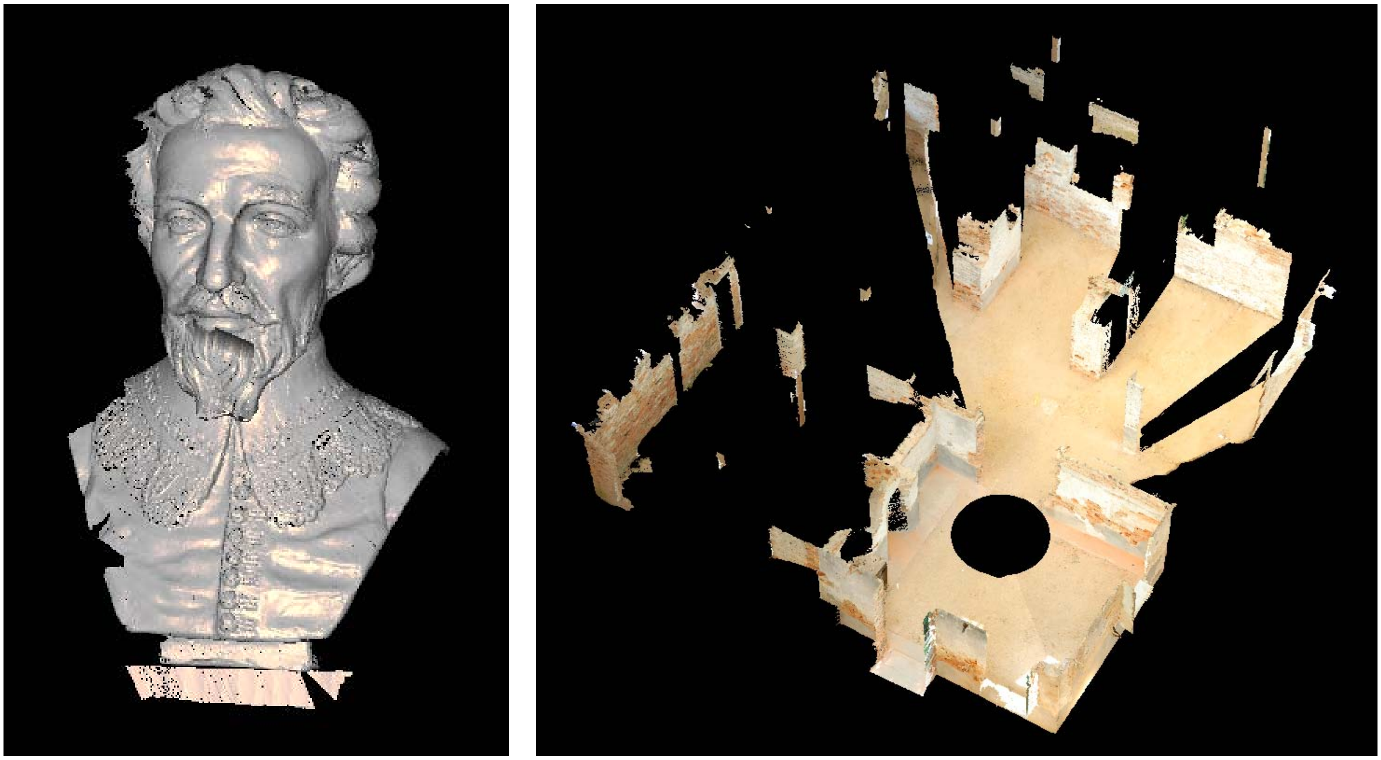

Figure 3: 3D scans of Raleigh Tavern bust (left) and Coffeehouse foundations (right) from Colonial Williamsburg datasets. The Coffeehouse scan contains 41 million points out of 2 billion total; color data is not maintained through the surface reconstruction phase of our processing pipeline.

2. Supercomputing for geometric analysis of scanned 3D models. Producing complete 3D digital models of cultural heritage artifacts then allows us to apply various algorithmic geometric analyses that can lead to new scholarly results. For example, in prior work, we have developed computer algorithms for automating the re-assembly of fragmented archaeological artifacts.[8] If collections of fragmented artifacts can be scanned, then these computer-aided "3D jigsaw puzzle solving" methods can greatly expedite many archaeological reconstruction problems.

The most effective computational geometry algorithms for automated fragment re-assembly, however, are very computationally expensive, and thus have traditionally been regarded as inefficient or useless for large reconstruction problems (i.e. re-assembling hundreds of fragments or more). However, once again, high performance supercomputers have the potential to overcome these limitations, if the algorithms and their software implementations can be adapted to execute in a massively parallel computing environment. As part of our project to investigate supercomputing applications for humanities scholarship, then, we have adapted our fragment re-

[8] D. Koller, "Protected Dissemination and Computer-aided Reconstruction of Digitized Cultural Artifacts", Ph.D. Dissertation, Stanford University, 2007.

assembly software to run in a high performance computing environment. This was a simple process, as each pairwise fragment-to-fragment matching candidate evaluation was considered as a separate computational task to be assigned to a compute node, and could execute independently of other tasks and thus no inter-process communication was necessary.

For evaluation purposes, we used the Stanford Digital Forma Urbis Romae dataset as input to the parallelized implementation. This dataset includes nearly 1200 large marble fragments from an ancient Roman map that were laser scanned at 0.25 millimeter resolution (Fig. 4). One of the parallelized matching algorithms is based upon finding correspondences between carved incisions on the top surfaces of the fragments. Another algorithm was implemented that samples the fragment fractured edges on finely spaced regular grids, and then performs brute force match evaluations across the set of sampled fragment edges. Such a brute force matching algorithm would be too costly on a uniprocessor computer, but a parallel implementation can speed up the run time by a factor of hundreds as compute nodes are available.

The output of our fragment matching algorithms is a list of scores for each potential pair of matching fragments, sorted to identify the most likely positive matches. Determining whether a proposed match is actually a new found mating of Forma Urbis fragments must be verified manually by an expert user familiar with the archaeology of the fragments, and the number of false positive matches proposed by the matching algorithms is extremely high (over 99% false positive rate). Thus, a large amount of time and effort is expended culling through incorrectly proposed matches. To date, we have manually examined 240 new possible candidate matches amongst the Forma Urbis fragments proposed by our parallel matching algorithms, but have unfortunately not been able to qualify any of them as actual new matches using the other available archaeological evidence (such as fragment color, quality of the marble, presence of cracks, etc.). However, we have many more proposed matches to verify, and are hopeful that the move to a high performance computing environment will allow us to discover some new fragment joins, which may even lead to some new discoveries about the topography of the ancient Roman city that is depicted on the map fragments.

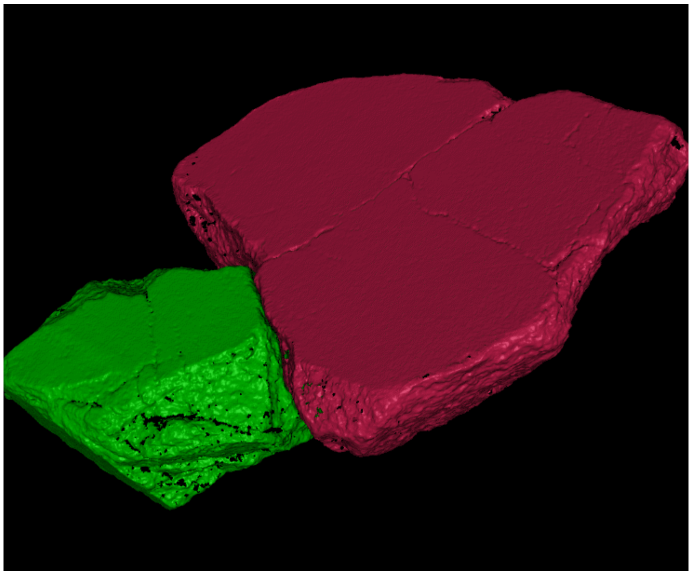

Figure 4. Computer matching algorithm searching for matches between the fractured edges of two digitized 3D fragments of the Forma Urbis Romae dataset.

Project Outcomes and Future Work

A core product of our research has been the codes for highly parallelized high performance computing implementations of algorithms for 3D scan registration, surface reconstruction, and Forma Urbis fragment matching.

However, although our codes are successful for processing cultural heritage scan data in massively parallel computing environments, we believe that they are quickly becoming matched in utility, and exceeded in usability, by widely available, commercial "off the shelf" 3D scan data processing software that runs on commodity computing hardware. In the recent past, mainstream commercial 3D scan data processing software that ran on typical desktop personal computers was not capable of processing 3D scanned models that consisted of hundreds of millions of sample points. However, even during the brief course of this project, 3D scan data processing software and affordable commodity workstations have continued to improve dramatically in performance and capacity such that these large models can be now processed on relatively inexpensive off-the-shelf systems. All the major scan processing packages (PolyWorks, RapidForm, Geomagic) are now available in 64-bit versions that can access many gigabytes of memory, and they support multithreading and out-of-core processing for datasets with billions of points. The commercial software does not yet perform the advanced, computationally intensive non-rigid registration algorithm that we used, but rigid registration is sufficient for most cultural heritage applications.

Because of the diminished value of our parallel 3D scan data processing codes relative to the rapidly evolving off-the-shelf systems, we anticipate little value in continuing this particular line of research. However, the algorithms and implementations for fragment reassembly are still fertile ground for new approaches and improvements.

In continuing work, we will investigate two other areas of high performance computing applied to 3D cultural heritage models. One area is remote interactive rendering of very large 3D models. Our "ScanView" remote rendering system[9] has been in use for 8 years providing protected access to statue models from the Stanford Digital Michelangelo Project, but utilizes a basic scanline renderer on a single GPU. High performance computing can allow for sophisticated ray-tracing and other rendering algorithms using highly parallel hardware, producing photorealistic renderings of huge 3D models at interactive rates for multiple users remotely accessing the rendering server. Secondly, we are interested in the application of HPC to perform accurate physical simulations based on 3D cultural heritage models. For example, with 3D models of historical structures and cities, we can simulate the effects of historical fires, floods, traffic flows, and other physical phenomena that would add insight to scholarly analysis and interpretation. We have been discussing these possible physical simulation applications with colleague Eric Field of the University of Virginia School of Architecture.

[9] D. Koller *et al.*, "Protected Interactive 3D Graphics Via Remote Rendering", *Proc. of ACM SIGGRAPH 2004*.

Our work on this project has given us a broad exposure to the use of high performance computing in the humanities, and we are preparing a manuscript that surveys prior existing humanities HPC applications, including details of our own work on applying HPC to digitized 3D cultural heritage models. We anticipate publication both in the computing research literature as well as a digital humanities venue, and believe this will contribute to efforts to make both communities of computer scientists and humanists aware of potential HPC collaborations.

Lessons Learned

Here we enumerate several general lessons learned in the course of this project, that may be useful to other researchers:

- There are many very powerful supercomputer resources available to humanities researchers. The National Energy Research Scientific Computing Center (NERSC), the National Science Foundation (NSF) TeraGrid, and the IBM World Community Grid all have encouraged humanities scholars to apply for allocations of HPC resources. Humanities applications can easily stand out from the traditional scientific and engineering computing tasks, and the supercomputing centers will champion your project!
- In our experience with NERSC, their staff is extremely helpful and the system documentation and account administration tools are excellent.
- There are many computer scientists (including the Project Director) who appreciate the opportunity to work with humanities researchers, and will gladly assist in applying computational methods to humanities applications.
- Converting pre-existing conventional software codes for usage on supercomputing systems took much longer than expected, due to different operating systems, software library versions, and job control environments.
- The overhead of transferring large datasets to/from a remote supercomputing center can nullify some or much of the advantages of remote HPC resources (we experienced ~1 MB/sec sustained transfer rate). Our very large 3D cultural heritage datasets, large both in input and output, were thus troublesome, as we wanted to frequently download processed 3D models for visualization. Local HPC systems at universities may thus be more convenient than the supercomputing sites.
- Commercial off-the-shelf 3D scan processing software and commodity multiprocessor workstations can now handle the processing of scanned datasets containing over one billion points into a finished 3D model. This capacity should be suitable for the vast majority of humanities and cultural heritage 3D models, and supercomputers will not be necessary.

Acknowledgements

Portions of the software code were provided by Benedict Brown, Brian Curless, the Stanford Graphics Laboratory, Michael Booth, and Tyler Eldred. Experimental 3D scan data was provided by the Stanford Digital Michelangelo Project, the University of Virginia Institute for Advanced Technology in the Humanities, and the Colonial Williamsburg Foundation. The Forma Urbis data is courtesy of the Stanford Digital Forma Urbis Romae Project.